

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Utility Patent Application

GAZE MANIPULATION

Inventor(s):

Antonio Criminisi

Andrew Blake

Philip Torr

Jamie Shotton

CLIENT'S DOCKET NO. MS304561.1

ATTORNEY'S DOCKET NO. MS1-1631US

EL996276579

1
2 GAZE MANIPULATION
3

4 Technical Field
5

6 The invention relates generally to image processing, and more particularly
7 to gaze manipulation using image processing.
8

9 Description
10

11 Digital video cameras are useful in both consumer and professional
12 contexts. Generally, digital video cameras capture sequences of digital images,
13 which may then be transferred to a computer system for display or processing or to
14 a storage device for storage.
15

16 One specific practice employs a digital video camera in a video
17 conferencing application. In a typical video conference, an image sequence
18 depicting a conference participant is transmitted to one or more other participants.
19 Concurrently, image sequences depicting the other participants are transmitted to
20 the first participant's display device. In this manner, each participant can view an
21 interactive video of the other participants during the conference.
22

23 In a typical video teleconferencing environment, a single video camera is
24 focused on a conference participant, who views the other participants in a video
25 window in his or her display device. The video camera is commonly mounted on
or near the display of a computer or television system in an attempt to minimize
the angle between the camera and the video window. Minimizing this angle can
enhance the illusion that the participant is looking into the camera instead of at the
video window in the display device. However, the angle is never really eliminated
and in fact can remain significant, particularly when the camera is close to the
26
27

1 participant. As a result, for example, the other participants receive a display of the
2 top or side of the first participant's head, instead of a straight-on view of the first
3 participant's face.

4 This situation provides a diminished user experience and limits the
5 effectiveness of such video conferencing. It is difficult to develop trust between
6 participants in the conference because of the difficulty in establishing eye contact
7 (i.e., the displayed participant is looking at his or her display instead of the
8 camera). Likewise, facial expressions may be distorted or obscured by the angular
9 discrepancy, thereby losing some important communication cues.

10 Conceptually, these problems may be resolved by a physical camera
11 positioned in the center of the display window, such that the participant's gaze and
12 the camera's axis are aligned - envision a video display with a hole drilled in the
13 middle of it in which to mount the camera. However, such configurations are
14 impractical for obvious reasons.

15 Some of these problems have been addressed by generating a cyclopean
16 virtual image, which approximates an image "captured" by a virtual camera
17 positioned between stereo cameras, whether centered or not-centered. Such
18 approaches, however, have not directly or sufficiently addressed certain artifacts
19 that result from attempting to combine a pair of stereo images. Some of these
20 artifacts are, for example, due to the difficulty of matching corresponding pixels
21 across images and reconstructing occluded regions.

22 For example, visualize a subject sitting in view of two horizontally
23 positioned stereo cameras. One camera may be able to view a doorknob behind
24 and slightly to the side of the subject's head, while the second camera cannot view
25 the doorknob because the subject's head occludes the doorknob from the second

1 camera's view. In the previously mentioned cyclopean virtual imaging
2 techniques, and in other approaches, the difficulty of reconstructing such occluded
3 regions in the cyclopean image often results in a "halo" effect around the subject's
4 head as well as other artifacts in the image.

5 Implementations described and claimed herein address these problems with
6 a multi-layer graph for dense stereo dynamic programming to improve processing
7 of occluded regions. Use of a multi-layer graph for dense stereo dynamic
8 programming allows an algorithm to distinguish between stereo disparities caused
9 by occlusion and disparities caused by non-fronto-parallel surfaces. This
10 distinction can be leveraged to reduce the occlusion effects, such as "halos". In
11 addition, anisotropic smoothing may be used to improve scan line matching. Also,
12 temporal background layer maintenance may be used to reduce temporal artifacts.

13 In various implementations, articles of manufacture are provided as
14 computer program products. One implementation of a computer program product
15 provides a computer program storage medium readable by a computer system and
16 encoding a computer program. Another implementation of a computer program
17 product may be provided in a computer data signal embodied in a carrier wave by
18 a computing system and encoding the computer program. The computer program
19 product encodes a computer program for executing on a computer system a
20 computer process that computes a minimum cost path in a stereo disparity model
21 between a scan line of a first image and a corresponding scan line of a second
22 image of a stereo image pair. The stereo disparity model distinguishes between
23 non-fronto-parallel matched pixels in each scan line and occluded pixels in each
24 scan line.

25

1 In another implementation, a method computes a minimum cost path in a
2 stereo disparity model between a scan line of a first image and a corresponding
3 scan line of a second image of a stereo image pair. The stereo disparity model
4 distinguishes between non-fronto-parallel matched pixels in each scan line and
5 occluded pixels in each scan line.

6 In yet another implementation, a system provides a dynamic programming
7 module computing a minimum cost path in a stereo disparity model between a
8 scan line of a first image and a corresponding scan line of a second image of a
9 stereo image pair. The stereo disparity model distinguishes between non-fronto-
10 parallel matched pixels in each scan line and occluded pixels in each scan line.

11 Other implementations are also described and recited herein.

12 Brief descriptions of the drawings included herein are listed below.

13 FIG. 1 illustrates an exemplary system for generating a cyclopean virtual
14 image with gaze manipulation.

15 FIG. 2 illustrates an exemplary video conferencing system
16 configuration 200 for generating a cyclopean virtual image with gaze
17 manipulation.

18 FIG. 3 illustrates stereo disparity as a function of left and right epipolar
19 lines L and R, which are defined in terms of pixel coordinates m and n ,
20 respectively.

21 FIG. 4 illustrates disparity and cyclopean axes overlaid on the L and R
22 axes.

23 FIG. 5 illustrates an exemplary five-move disparity process model.

24 FIGS. 6, 7, and 8 combine to represent an exemplary three-plane
25 representation of the five-move disparity model.

1 FIG. 9 illustrates an exemplary stereo disparity graph for matched points.

2 FIG. 10 illustrates an exemplary stereo disparity graph for occluded points.

3 FIG. 11 illustrates exemplary operations for performing gaze manipulation.

4 FIG. 12 illustrates an exemplary four-move disparity process model.

5 FIGs. 13, 14, and 15 combine to represent a four-plane representation of the
6 four-move disparity model.

7 FIG. 16 illustrates an exemplary system useful for implementing an
8 embodiment of the present invention.

9 A multi-layer graph for dense stereo dynamic programming can improve
10 synthesis of cyclopean virtual images by distinguishing between stereo disparities
11 causes by occlusion and disparities caused by non-fronto-parallel surfaces. This
12 distinction can be leveraged to reduce the occlusion effects, such as “halos”.
13 Generally, occlusion refers to an image region that is captured by one camera but
14 is not captured by a second camera because of an obstruction in the second
15 camera’s view (e.g., a video conference participant’s head occluding a region of
16 the background from the second camera’s view). Furthermore, a fronto-parallel
17 surface refers to a surface that is substantially parallel to an axis connecting the
18 left and right cameras. The axis may be referred to as the “baseline”. As such, a
19 non-fronto-parallel surface slants away from one camera or the other.
20 Distinguishing at least between these two types of disparities allows improved
21 matching of left and right pixel data, which increases the amount of correct pixel
22 information used in constructing the cyclopean virtual image and minimizes
23 occlusion artifacts.

24 FIG. 1 illustrates an exemplary system 100 for generating a cyclopean
25 virtual image with gaze manipulation. In the system 100, a left image 102 is

1 captured by a camera mounted on the right side of the video display, as seen by
2 the user. Likewise, a right image 104 is captured by a camera mounted on the left
3 side of the video display, as seen by the user. As such, in both images, the user
4 can be seen looking into the video display, as opposed to looking directly at one of
5 the cameras. The left and right images 102 and 104 are input to a dynamic
6 programming module 106, which generates a stereo disparity graph for each
7 corresponding pair of epipolar lines of the images 102 and 104. In the illustrated
8 implementation, a three-plane model for the dynamic programming is used,
9 although other graphs may be employed, such as a four-plane model, etc.

10 The stereo disparity graph generated by the dynamic programming module
11 is input to a cyclopean virtual image generator 108, which uses pixel
12 characteristics of corresponding pixels associated with a minimum cost path in the
13 stereo disparity graph to generate the cyclopean virtual image 110 with gaze
14 correction. As a result, the cyclopean virtual image 110 shows the user as
15 appearing to look directly into the camera.

16 FIG. 2 illustrates an exemplary video conferencing system
17 configuration 200 for generating a cyclopean virtual image with gaze
18 manipulation. A computer system 202 is coupled to a video display 204 having
19 two cameras 206 and 208 mounted on either side of the video display 204. A
20 video window 210 displays a remote participant on the other end of the video
21 conference session.

22 In a configuration having only a single camera, the user typically focus his
23 or her eyes on the video window 210, while the single camera captures images of
24 the user from one side of the other. As such, the captured images sent to the
25 remote participant are primarily a side view of the user's head, not a straight-on

1 view of the user's face. The illustrated configuration, however, allows generation
2 of a cyclopean virtual image from the captured left and right images of the user.

3 The cyclopean virtual image generation synthesizes the cyclopean virtual
4 image from a stereo disparity graph representing the disparity field between
5 corresponding left and right images. Furthermore, the dynamic programming
6 applied to the disparity graph distinguishes between disparities caused by
7 occlusion and disparities caused by non-fronto-parallel surfaces in the view field.

8 It should be understood that more than two cameras may also be used to
9 generate a cyclopean virtual image. Likewise, the cameras may be in alternative
10 orientations, such as at the top and bottom of the video display. For example, one
11 configuration may include four cameras, each placed at a corner of the video
12 display.

13 The cyclopean virtual image \hat{I} is synthesized from intensity functions
14 $L = \{L_m, m = 0, \dots, N\}$ and $R = \{R_n, n = 0, \dots, N\}$, which represent epipolar lines (or
15 scan lines) of observed (i.e., captured) left and right images. A matched pair $(L_m,$
16 $R_n)$ has "stereo disparity" of $d=n-m$, which may be considered a measure of
17 "parallax". In one implementation, each image contains color pixels in three color
18 channels, such that $L_m, R_n \in \Re^3$. In a more general setting, however, there may be
19 other features, such that $L_m, R_n \in \Re^f$, where f is an integer. For example, groups of
20 pixels may be filtered to obtain improved invariance to illumination variations or
21 non-uniform camera sensitivities.

22 A cyclopean epipolar line (i.e., the corresponding scan line in the virtual
23 cyclopean image) is represented by $I = \{I_k, k = 0, \dots, 2N\}$. The cyclopean virtual
24 image \hat{I} is constructed from a set of cyclopean epipolar lines stacked line-by-line
25 to form the resulting cyclopean image.

1 FIG. 3 illustrates stereo disparity as a function of left and right epipolar
2 lines **L** and **R**, which are defined in terms of pixel coordinates m and n ,
3 respectively. The stereo disparity between the left and right stereo images is
4 defined as a vector $\mathbf{d} = \{d_k, k = 0, \dots, 2N\}$ having components expressed in cyclopean
5 coordinates k .

6 A diagram 300 shows an axis 302, representing a sequence of positions
7 along a left epipolar line **L**, and another axis 304 representing a sequence of
8 positions along a right epipolar line **R**. A minimum cost path 306 indicates
9 matches between pixels in given sequential positions in **L** with pixels in given
10 sequential positions in **R**. For example, pixel 1 of **L** matches pixel 1 of **R**, as
11 shown by point 308. In contrast, pixel 3 of **L** matches pixel 2 of **R**, as shown by
12 point 310. The disparity associated with a point 310 on the minimum cost
13 path 306 is defined as the orthogonal distance of the point from a virtual scan
14 line 312 (or zero disparity axis or zero parallax axis). For example, the disparity
15 of the point 308 is zero, whereas the disparity d of the point 310 is shown by
16 line 314. (As suggested by the disparity axis of FIG. 4, the disparity of point 310
17 is “-1”.)

18 Accordingly, the minimum cost path 306 represents a two-dimensional
19 profile of a scan line of the virtual image, where pixels with a greater absolute
20 value of disparity (e.g., point 310, which has a negative disparity relative to the
21 zero parallax line 312) are closer to the virtual cyclopean camera – e.g., the video
22 subject - than pixels with a lower absolute value of disparity (e.g., point 316,
23 which has a zero disparity relative to the zero parallax line 312), which are
24 deemed farther away from the virtual cyclopean camera – e.g., the background.
25 Stacking a set of these two-dimensional profiles, which correspond to individual

1 cyclopean epipolar lines, can yield a three-dimensional profile surface of the
2 image subject.

3 A matching cost function may be used to determine the minimum cost path
4 in a stereo disparity graph. A variety of matching cost functions may be employed
5 to compute the matching two pixels. However, using some traditional techniques,
6 processing individual epipolar line pairs independently can cause visible “streaky”
7 artifacts in the output disparity graph. Therefore, by using neighborhood windows
8 in computing the cost of matching two pixels, the “independence” of the scan lines
9 can be compromised, thereby reducing streaky artifacts.

10 In one implementation, a windowed Normalized Sum of Squared
11 Differences (SSD) matching function is used to compute the matching cost $M(l, r)$
12 for every pair of pixels along corresponding epipolar lines:

$$13 M_{ssd}(l, r) = \frac{M'(l, r)}{2} \quad (1)$$

15 with

$$16 M'(l, r) = \frac{\sum_{\delta \in \Omega} [(I'_{p_l + \delta} - \bar{I}'_{p_l}) - (I'_{p_r + \delta} - \bar{I}'_{p_r})]^2}{\sum_{\delta \in \Omega} (I'_{p_l + \delta} - \bar{I}'_{p_l})^2 + \sum_{\delta \in \Omega} (I'_{p_r + \delta} - \bar{I}'_{p_r})^2} \quad (2)$$

19 where Ω is an $n \times m$ generic template patch centered at the origin of the coordinate
20 system; p_l and p_r are pixel positions (2-vectors) in the left and right images,
21 respectively; and δ is a variable 2D displacement vector. The “bar” above a
22 variable (e.g., \bar{I}) represents the mean operator.

23 In FIGs. 6, 7, and 8 combine to represent a 3-plane representation of the 5-
24 move disparity model another implementation, a Normalized Cross-Correlation
25 (NCC) matching cost may be employed:

$$M_{ncc}(l, r) = \frac{1 - M'(l, r)}{2} \quad (3)$$

where

$$M'(l, r) = \frac{\sum_{\delta \in \Omega} (I'_{p_l+\delta} - \bar{I}'_{p_l})(I'_{p_r+\delta} - \bar{I}'_{p_r})}{\sqrt{\sum_{\delta \in \Omega} (I'_{p_l+\delta} - \bar{I}'_{p_l})^2} \sum_{\delta \in \Omega} (I'_{p_r+\delta} - \bar{I}'_{p_r})^2} \quad (4)$$

is the correlation coefficient. Other matching cost functions may also be used, including without limitation shiftable window approaches (e.g., using 3×3 pixel windows or larger) or rectangular window approaches (e.g., using 3×7 windows).

FIG. 4 illustrates disparity and cyclopean axes overlaid on the **L** and **R** axes to show an exemplary stereo disparity graph 400. Based on the disparity axis 402, a disparity vector **d** in cyclopean coordinates **k** along the cyclopean axis 404 can be graphed into the pixel coordinates *m* and *n*. The cyclopean coordinate *k* corresponding to pixel coordinates *m* and *n* is computed as $k=m+n$. The bold line marks the minimum cost path 406 in the stereo disparity graph 400.

Different segments of the minimum cost path 406 represent different characteristics of the stereo images. A diagonal path on the $d=0$ axis (as seen between $k=0$ to 2) represents a zero-disparity, linear match between pixels in the epipolar lines of the right and left images. This linear match might happen, for example, when the pixels are of distant objects in which no parallax is evident. In contrast, a diagonal path off of the $d=0$ axis (as seen between $k=3$ to 5) represents a disparate (disparity = -1), linear match between pixels in the epipolar lines of the right and left images. In both cases, a diagonal line on the minimum cost path 406 represents matched pixels.

Horizontal and vertical lines (as seen between $d=2$ to 3) in the minimum cost path 406 have traditionally been considered to represent only occluded

1 regions. For example, in FIG. 4, horizontal lines would be deemed to indicate
2 pixels that are occluded from the right camera, while vertical lines would be
3 deemed to indicate pixels that are occluded from the left camera.

4 However, in an approach described herein, horizontal and vertical lines are
5 considered to indicate at least either occluded pixels or matched pixels of non-
6 fronto-parallel surfaces. Non-fronto-parallel surfaces cause multiple pixels from
7 one camera image to match with a single pixel in the other camera image, thereby
8 inducing a horizontal or vertical line in the stereo disparity graph.

9 FIG. 5 illustrates an exemplary 5-move disparity process model 500. The
10 points 502, 504, 506, and 508 represent possible pixels in the stereo disparity
11 graph, such that the diagonal axis 510 represents a diagonal move 512 between
12 pixels in a stereo disparity graph. The horizontal axis 514 represents a horizontal
13 move between pixels in a stereo disparity graph and the vertical axis 516
14 represents a horizontal move between pixels in a stereo disparity graph.

15 However, as discussed above, horizontal and vertical moves (i.e., non-
16 diagonal moves) can represent at least either occluded pixels or matched pixels of
17 non-fronto-parallel surfaces. Therefore, two categories of such moves are
18 designated in each direction: (non-fronto-parallel) matched moves (518 and 520)
19 and occluded moves (522 and 524). As such, FIG. 5 illustrates a 5-move disparity
20 process model, although a 4 move model may also be employed.

21 FIGs. 6, 7, and 8 combine to represent a 3-plane representation of the 5-
22 move disparity model, but they are split out into separate figures for clarity. The 5
23 move model applies to moves between adjacent pixels in the stereo disparity
24 graph. In one implementation, to distinguish between (non-fronto-parallel)

25

1 matched moves and occluded moves, three planes are used: a left-occluded plane
2 L, a matched plane M, and a right-occluded plane R.

3 In FIG. 6, the moves from an occluded plane to the matched plane are
4 shown (from empty circle to filled circle) in model portion 600. A cost penalty of
5 β is applied to these moves. In FIG. 7, the moves 700 and 702 from the matched
6 plane to an occluded plane are shown (from empty circle to filled circle) in model
7 portion 704. A cost penalty of β is applied to these moves. Also in FIG. 7, the
8 moves 706 and 708 from one pixel in an occluded plane to another pixel in the
9 same occluded plane are shown (from empty circle to filled circle) in the model
10 portion 704. A cost penalty of α is applied to these moves. In one
11 implementation, α is set to 0.5 and β is set to 1.0, although other value
12 combinations are also contemplated. In FIG. 8, the moves from one pixel in the
13 matched plane to another pixel in the matched plane are shown (from empty circle
14 to filled circle) in the model portion 800. No cost penalty is applied to these
15 moves.

16 The 3-plane model provides a basis for altering the individual costs to
17 distinguish between different types of moves. For example, biasing the penalty
18 costs against inter-plane moves tends to keep runs of occluded or non-occluded
19 pixels together, thus reducing most of the inaccuracies in the reconstruction of
20 occlusions and disparities. Also, logically impossible moves, such as the direct
21 transition between left and right occlusions are prohibited simply by removing
22 certain transitions from the set of allowed transitions in the 3-plane graph.

23 In one implementation, the cost $C(A \rightarrow B)$ of a generic transition between
24 two planes A and B is manually set, but it is also possible to set $C(A \rightarrow B)$
25 probabilistically. Moreover, it may be assumed that $C(A \rightarrow B)$ is symmetric (i.e.,

1 $C(A \rightarrow B) = C(B \rightarrow A)$). This assumption leads to the two penalty parameters: α
 2 being the penalty for a move within an occluded plane, and β being the cost of a
 3 move between different planes.

4 As such, in this exemplary implementation, the matrices of cumulative
 5 costs C_L , C_M , and C_R (one for each plane in the graph) are initialized to ∞
 6 everywhere except in the right occluded plane, where:

$$7 \quad C_R(i,0) = i\alpha \quad (5)$$

9 and the forward step of the dynamic programming proceeds as follows:

$$11 \quad C_L(l,r) = \min \begin{cases} C_L(l,r-1) + \alpha \\ C_M(l,r-1) + \beta \end{cases} \quad (6)$$

$$13 \quad C_M(l,r) = M(l,r) + \min \begin{cases} C_M(l-1,r) \\ C_L(l-1,r) + \beta \\ C_R(l-1,r) + \beta \\ C_M(l,r-1) \end{cases}$$

$$16 \quad C_M(l,r) = M(l,r) + \min \begin{cases} C_L(l,r-1) + \beta \\ C_R(l,r-1) + \beta \\ C_M(l-1,r-1) \\ C_L(l-1,r-1) + \beta \\ C_R(l-1,r-1) + \beta \end{cases} \quad (7)$$

$$20 \quad C_R(l,r) = \min \begin{cases} C_R(l-1,r) + \alpha \\ C_M(l-1,r) + \beta \end{cases} \quad (8)$$

22 wherein $M(l,r)$ is the cost of matching the l^{th} pixel in the left scan line with the r^{th}
 23 pixel in the right scan line.

24 Based on these costs, the minimum cost path is determined for the scan line
 25 pair. The matching cost computation and the dynamic programming are repeated

1 for each scan line pair in the stereo images. The synthesis of the cyclopean virtual
2 view can be done for each scan line by taking a point \mathbf{p} on the minimum cost path,
3 taking the colors of the corresponding pixels \mathbf{p}_l and \mathbf{p}_r in the left and right scan
4 lines, averaging them together, and projecting the newly obtained pixel
5 orthogonally to the virtual image plane into the virtual image point \mathbf{p}_v .

6 FIG. 9 illustrates an exemplary stereo disparity graph for matched points.
7 A stereo disparity graph 900 shows an axis 902, representing a sequence of
8 positions along a left scan line \mathbf{L} , and another axis 904 representing a sequence of
9 positions along a right scan line \mathbf{R} . The minimum cost path 906 indicates
10 minimum cost matches between pixels in given sequential positions in \mathbf{L} with
11 pixels in given sequential positions in \mathbf{R} . The disparity associated with a point on
12 the minimum cost path 906 is defined as the orthogonal distance of the point from
13 a virtual scan line 908.

14 A matched point \mathbf{p} is projected orthogonally onto its corresponding point \mathbf{p}_v
15 on the virtual scan line 908 to designate the position of the corresponding
16 cyclopean virtual image pixel on the virtual scan line. The pixel value of the
17 virtual pixel \mathbf{p}_v is the average of the corresponding pixels \mathbf{p}_l and \mathbf{p}_r .

18 FIG. 10 illustrates an exemplary stereo disparity graph for occluded points.
19 A stereo disparity graph 1000 shows an axis 1002, representing a sequence of
20 positions along a left scan line \mathbf{L} , and another axis 1004 representing a sequence of
21 positions along a right scan line \mathbf{R} . The minimum cost path 1006 indicates
22 minimum cost matches between pixels in given sequential positions in \mathbf{L} with
23 pixels in given sequential positions in \mathbf{R} . The disparity associated with a point on
24 the minimum cost path 1006 is defined as the orthogonal distance of the point from
25 a virtual scan line 1008.

1 An occluded point \mathbf{p} on the continuation 1010 of the background (with the
2 same disparity) is projected orthogonally onto its corresponding point \mathbf{p}_v on the
3 virtual scan line 1008. Because \mathbf{p} represents a pixel within a left occlusion in this
4 illustration, the pixel value of \mathbf{p}_v is the same as that of the corresponding point \mathbf{p}_r
5 on the right view only.

6 FIG. 11 illustrates exemplary operations for performing gaze manipulation.
7 A receipt operation 1100 receives the stereo images from the stereo cameras. A
8 computation operation 1102 computes the matching cost for each pixel of the
9 epipolar line pair. A filtering operation 1104 filters the matching costs to reduce
10 streaky artifacts caused by scan line independence.

11 A dynamic programming operation 1106 alters the initially computed
12 individual costs for each pixel pair to designate different types of moves and
13 therefore different types of disparities (e.g., occlusion disparities versus non-
14 fronto-parallel disparities). Based on the altered costs, a minimum cost path is
15 identified in a path operation 1108. An imaging operation 1110 determines the
16 cyclopean virtual scan line based on the minimum cost path in the stereo disparity
17 graph.

18 While FIGs. 5-8 address a five-move, three-plane disparity model, other
19 models may also be employed. For example, a four-move, four-plane model can
20 prove as reliable and easier to use. In particular, in one implementation of a four-
21 move mode, every possible path through the cost space has the same length (i.e.,
22 the same Manhattan distance between the opposite corners of the cost space), thus
23 making the comparison of path costs more meaningful. Furthermore, the removal
24 of the diagonal move (see move 512 in FIG. 5) makes the model symmetrical and
25 thus more suitable for a possible probabilistic formulation.

1 FIG. 12 illustrates an exemplary four-move disparity process model 1200.
2 The points 1202, 1204, 1206, and 1208 represent possible pixels in the stereo
3 disparity graph, such that the diagonal axis 1210 a zero-disparity axis in a stereo
4 disparity graph. However, in the four-move model 1200, no diagonal move is
5 modeled. The horizontal axis 1214 represents a horizontal move between pixels in
6 a stereo disparity graph and the vertical axis 1216 represents a horizontal move
7 between pixels in a stereo disparity graph.

8 As discussed above, horizontal and vertical moves (i.e., non-diagonal
9 moves) can represent at least either occluded pixels or matched pixels of non-
10 fronto-parallel surfaces. Therefore, two categories of such moves are designated
11 in each direction: (non-fronto-parallel) matched moves (1218 and 1220) and
12 occluded moves (1222 and 1224). As such, FIG. 12 illustrates a four-move
13 disparity process model.

14 FIGs. 13, 14, and 15 combine to represent a four-plane representation of the
15 four-move disparity model, but they are split out into separate figures for clarity.
16 The four-move model applies to moves between adjacent pixels in the stereo
17 disparity graph. In the illustrated implementation, to distinguish between (non-
18 fronto-parallel) matched moves and occluded moves, four planes are used: a left-
19 occluded plane L_O , a left matched plane L_M , a right matched plane R_M , and a right-
20 occluded plane R_O . In this model, a typical “matched” move, which in a five-
21 move model would involve a diagonal move, would involve two matched moves,
22 one vertical and one horizontal in a 2D graph or a two-move oscillation between
23 two adjacent matched planes (e.g., from L_M to R_M and back to L_M).

24 In FIG. 13, the moves within individual planes are shown (from empty
25 circle to filled circle) in model portion 1300. Again, logically impossible moves,

such as the direct transition between left and right occlusions are prohibited simply by removing certain transitions from the set of allowed transitions in the four-plane graph. A cost penalty of α is applied to the moves with the occluded planes L_O and R_O , and a cost penalty of $\gamma+M(l,r)$ for moves within the matched planes L_M and R_M . In FIG. 14, the moves between an occluded plane and an adjacent matched plane are shown (from empty circle to filled circle) in model portion 1400. A cost penalty of $\beta+M(l,r)$ is applied to moves from an occluded plane to an adjacent matched plane, a cost penalty of β is applied to moves from a matched plane to an adjacent occluded plane, and a cost penalty of $M(l,r)$ is applied to moves between matched planes. In FIG. 15, the moves between an occluded plane and a non-adjacent matched plane are shown (from empty circle to filled circle) in model portion 1500. A cost penalty of $\beta+M(l,r)$ is applied to moves from an occluded plane to a non-adjacent matched plane, and a cost penalty of β is applied to moves from a matched plane to a non-adjacent occluded plane. In one implementation, α is set to 0.5, β is set to 1.0, and γ is set to 0.25, although other value combinations are also contemplated.

The four-plane model provides a basis for altering the individual costs to distinguish between different types of moves. For example, heavily biasing the penalty costs against moves in and out of an occluded plane tends to keep runs of occluded pixels together, thus reducing most of the inaccuracies in the reconstruction of occlusions and disparities. Therefore, once a path enters an occluded plane, the path is encouraged to stay in that plane unless a pair of strongly matched pixels is found (i.e., low $M(l,r)$ cost). In contrast, biasing moves within a single matched plane, albeit less heavily, discourages runs of matched moves, thereby favoring surfaces that are close to fronto-parallel. Hence, in this

model, slanted surfaces are modeled as oscillations between the two matched planes.

As such, in this exemplary implementation, the matrices of cumulative costs $C_{L_o}, C_{L_m}, C_{R_m}$, and C_{R_o} (one for each plane in the graph) are initialized to $+\infty$ everywhere except in the right occluded plane, where:

$$C_{R_o}(i,0) = i\alpha \quad (9)$$

and the forward step of the dynamic programming proceeds as follows:

$$C_{L_o}(l,r) = \min \begin{cases} C_{L_o}(l,r-1) + \alpha \\ C_{L_m}(l,r-1) + \beta \\ C_{R_m}(l,r-1) + \beta \end{cases} \quad (9)$$

$$C_{L_m}(l,r) = M(l,r) + \min \begin{cases} C_{L_m}(l,r-1) + \gamma \\ C_{R_m}(l,r-1) \\ C_{L_o}(l,r-1) + \beta \\ C_{R_o}(l,r-1) + \beta \end{cases} \quad (10)$$

$$C_{R_m}(l,r) = M(l,r) + \min \begin{cases} C_{L_m}(l-1,r) \\ C_{R_m}(l-1,r) + \gamma \\ C_{L_o}(l-1,r) + \beta \\ C_{R_o}(l-1,r) + \beta \end{cases} \quad (11)$$

$$C_{R_o}(l,r) = \min \begin{cases} C_{R_o}(l-1,r) + \alpha \\ C_{L_m}(l-1,r) + \beta \\ C_{R_m}(l-1,r) + \beta \end{cases} \quad (12)$$

wherein $M(l,r)$ is the cost of matching the l^{th} pixel in the left scan line with the r^{th} pixel in the right scan line.

Based on these costs, the minimum cost path is determined for the scan line pair. The matching cost computation and the dynamic programming are repeated

1 for each scan line pair in the stereo images. The synthesis of the cyclopean virtual
2 view can be done for each scan line by taking a point p on the minimum cost path,
3 taking the colors of the corresponding pixels p_l and p_r in the left and right scan
4 lines, averaging them together, and projecting the newly obtained pixel
5 orthogonally to the virtual image plane into the virtual image point p_v .

6 The exemplary hardware and operating environment of FIG. 16 for
7 implementing the invention includes a general purpose computing device in the
8 form of a computer 20, including a processing unit 21, a system memory 22, and a
9 system bus 23 that operatively couples various system components include the
10 system memory to the processing unit 21. There may be only one or there may be
11 more than one processing unit 21, such that the processor of computer 20
12 comprises a single central-processing unit (CPU), or a plurality of processing
13 units, commonly referred to as a parallel processing environment. The computer
14 20 may be a conventional computer, a distributed computer, or any other type of
15 computer; the invention is not so limited.

16 The system bus 23 may be any of several types of bus structures including a
17 memory bus or memory controller, a peripheral bus, and a local bus using any of a
18 variety of bus architectures. The system memory may also be referred to as
19 simply the memory, and includes read only memory (ROM) 24 and random access
20 memory (RAM) 25. A basic input/output system (BIOS) 26, containing the basic
21 routines that help to transfer information between elements within the computer
22 20, such as during start-up, is stored in ROM 24. The computer 20 further
23 includes a hard disk drive 27 for reading from and writing to a hard disk, not
24 shown, a magnetic disk drive 28 for reading from or writing to a removable
25

1 magnetic disk 29, and an optical disk drive 30 for reading from or writing to a
2 removable optical disk 31 such as a CD ROM or other optical media.

3 The hard disk drive 27, magnetic disk drive 28, and optical disk drive 30
4 are connected to the system bus 23 by a hard disk drive interface 32, a magnetic
5 disk drive interface 33, and an optical disk drive interface 34, respectively. The
6 drives and their associated computer-readable media provide nonvolatile storage
7 of computer-readable instructions, data structures, program modules and other
8 data for the computer 20. It should be appreciated by those skilled in the art that
9 any type of computer-readable media which can store data that is accessible by a
10 computer, such as magnetic cassettes, flash memory cards, digital video disks,
11 Bernoulli cartridges, random access memories (RAMs), read only memories
12 (ROMs), and the like, may be used in the exemplary operating environment.

13 A number of program modules may be stored on the hard disk, magnetic
14 disk 29, optical disk 31, ROM 24, or RAM 25, including an operating system 35,
15 one or more application programs 36, other program modules 37, and program
16 data 38. A user may enter commands and information into the personal computer
17 20 through input devices such as a keyboard 40 and pointing device 42. Other
18 input devices (not shown) may include a microphone, joystick, game pad, satellite
19 dish, scanner, or the like. These and other input devices are often connected to the
20 processing unit 21 through a serial port interface 46 that is coupled to the system
21 bus, but may be connected by other interfaces, such as a parallel port, game port,
22 or a universal serial bus (USB). A monitor 47 or other type of display device is
23 also connected to the system bus 23 via an interface, such as a video adapter 48.
24 In addition to the monitor, computers typically include other peripheral output
25 devices (not shown), such as speakers and printers.

1 The computer 20 may operate in a networked environment using logical
2 connections to one or more remote computers, such as remote computer 49. These
3 logical connections are achieved by a communication device coupled to or a part
4 of the computer 20; the invention is not limited to a particular type of
5 communications device. The remote computer 49 may be another computer, a
6 server, a router, a network PC, a client, a peer device or other common network
7 node, and typically includes many or all of the elements described above relative
8 to the computer 20, although only a memory storage device 50 has been illustrated
9 in FIG. 16. The logical connections depicted in FIG. 16 include a local-area
10 network (LAN) 51 and a wide-area network (WAN) 52. Such networking
11 environments are commonplace in office networks, enterprise-wide computer
12 networks, intranets and the Internet, which are all types of networks.

13 When used in a LAN-networking environment, the computer 20 is
14 connected to the local network 51 through a network interface or adapter 53,
15 which is one type of communications device. When used in a WAN-networking
16 environment, the computer 20 typically includes a modem 54, a type of
17 communications device, or any other type of communications device for
18 establishing communications over the wide area network 52. The modem 54,
19 which may be internal or external, is connected to the system bus 23 via the serial
20 port interface 46. In a networked environment, program modules depicted relative
21 to the personal computer 20, or portions thereof, may be stored in the remote
22 memory storage device. It is appreciated that the network connections shown are
23 exemplary and other means of and communications devices for establishing a
24 communications link between the computers may be used.

25

1 In an exemplary implementation, a dynamic programming module, a
2 cyclopean virtual image generator, and other modules may be incorporated as part
3 of the operating system 35, application programs 36, or other program
4 modules 37. The stereo disparity graph data, matching costs, altered costs, and
5 cyclopean virtual image data may be stored as program data 38.

6 The embodiments of the invention described herein are implemented as
7 logical steps in one or more computer systems. The logical operations of the
8 present invention are implemented (1) as a sequence of processor-implemented
9 steps executing in one or more computer systems and (2) as interconnected
10 machine modules within one or more computer systems. The implementation is a
11 matter of choice, dependent on the performance requirements of the computer
12 system implementing the invention. Accordingly, the logical operations making
13 up the embodiments of the invention described herein are referred to variously as
14 operations, steps, objects, or modules.

15 The above specification, examples and data provide a complete description
16 of the structure and use of exemplary embodiments of the invention. Since many
17 embodiments of the invention can be made without departing from the spirit and
18 scope of the invention, the invention resides in the claims hereinafter appended.